Artificial Intelligence for Corporate Governance in the U.S. Defense Industrial Base: A Scholar-Practitioner Framework for Accountability, Ethics, and National Security

Satyadhar Joshi
Alumnus, International MBA, Bar-Ilan University, Israel
Alumnus, M.S. in Information Technology, Touro College, New York, USA
ORCID: 0009-0002-6011-5080
Email: satyadhar.joshi@gmail.com

***Abstract*—The U.S. defense industrial base confronts a confluence of governance challenges unprecedented in scope and consequence: financial fraud, systemic product quality failures, safety-critical software deficiencies, and supply chain opacity that collectively jeopardize national security. Watershed events—the Boeing 737 MAX disasters, persistent subcontracting irregularities, and recurrent Nunn-McCurdy cost breaches—illuminate the structural inadequacy of traditional governance mechanisms for an industry in which technical complexity, classified information flows, and sovereign risk intersect. This paper develops a comprehensive, theory-grounded framework for integrating Artificial Intelligence (AI) into the corporate governance architectures of U.S. defense contractors. Drawing on six canonical governance theories—agency, stewardship, resource dependency, stakeholder, institutional, and transaction cost economics—and a systematic review of the emerging AI-governance literature, we demonstrate how AI can operationalize ethical principles, substantially reduce information asymmetry, enable continuous auditing, and anticipate supply chain disruptions at scale. The proposed AI governance meta-layer is calibrated to the specific regulatory obligations imposed by the International Traffic in Arms Regulations (ITAR), the Defense Federal Acquisition Regulation Supplement (DFARS), and the Paperwork Reduction Act of 1995. Concurrently, this paper constitutes a formal public comment submitted in response to Federal Register Volume 91, Issue 103 (May 29, 2026), concerning the proposed extension of information collection requirement OMB Control Number 0704-0253 under DFARS Part 244. The analysis demonstrates how AI-enabled automation can enhance the quality, practical utility, and analytical rigor of contractor purchasing system data while materially reducing respondent burden—directly addressing the four areas of public comment invited by the Defense Acquisition Regulations System (DARS).**

***Index Terms*—Artificial Intelligence, Corporate Governance, Defense Industrial Base, National Security, Accountability, Agency Theory, Stakeholder Theory, Ethical AI, Risk Management, Sarbanes-Oxley, DFARS, ITAR, Subcontracting Oversight, Paperwork Reduction Act, Federal Register.**

I. Introduction: Governance Failure as a National Security Risk

Corporate governance failures in the defense sector are not merely matters of financial or reputational concern—they constitute direct threats to national security. When a missile guidance component is defective, when a satellite's attitude control software harbors an undetected edge-case fault, or when cost overruns on a major acquisition program are systematically concealed from congressional oversight, the consequences extend far beyond the balance sheet. The U.S. defense industrial base is unique among commercial sectors in that its governance failures carry sovereign-level risk: they can erode military readiness, embolden adversaries, and cost lives.

Against this backdrop, the emergence of agentic AI in 2025–2026—autonomous systems capable of reasoning, planning, and adapting across complex operational environments—presents a transformative opportunity for defense governance. AI offers the prospect of a *meta-governance layer*: a persistent, real-time oversight infrastructure that monitors financial integrity, manufacturing quality, regulatory compliance, and supply chain resilience simultaneously, surfacing risk signals that would otherwise remain invisible until a catastrophic failure event.

This paper advances three interrelated contributions. First, it synthesizes classical governance theory with the AI-governance literature to produce a unified framework applicable across the defense industrial base—from prime contractors to Tier-3 subcontractors. Second, it draws generalizable lessons from major corporate failures in defense-adjacent industries, demonstrating how AI governance mechanisms could have prevented or substantially mitigated each. Third, it applies this framework to the specific context of DFARS Part 244 contractor purchasing system oversight, with the dual purpose of advancing scholarship and serving as a formal public comment to DARS on the proposed extension of OMB Control Number 0704-0253.

The central thesis is as follows: classical governance frameworks remain necessary but are no longer sufficient for the complexity, pace, and stakes of modern defense contracting. The strategic integration of AI does not displace human judgment or board authority—it augments them, providing a data-verified, continuously updated foundation upon which sound governance decisions can be made. When embedded within robust institutional structures and guided by explicit ethical frameworks, AI can restore the defense industrial base's standing as a reliable steward of public resources and national security.

A. Structure of This Paper

Section II situates the analysis within the specific context of DFARS Part 244 and the Federal Register notice. Section III reviews the foundational principles of defense contractor governance. Section IV maps classical governance theories to the defense context. Section V extracts AI governance lessons from major corporate failures. Section VI maps AI capabilities to each governance theory. Section VII addresses the ethical imperative of preventing safety-critical failures. Section VIII provides targeted recommendations on AI-driven information collection improvement. Sections IX and X survey benefits and barriers. Section XI offers board-level strategic guidance. Section XII presents structured recommendations for the Federal Register response, and Section XIII draws broader policy implications.

B. Formal Response to Federal Register Notice DARS-2026-0202

This paper is submitted as a substantive public comment in response to the Federal Register notice published May 29, 2026 (Volume 91, Issue 103), in which the Defense Acquisition Regulations System announced the proposed extension of an information collection requirement under the Paperwork Reduction Act of 1995, OMB Control Number 0704-0253, covering DFARS Subcontracting Policies and Procedures (DFARS Part 244). The current OMB approval expires August 31, 2026; DoD proposes a three-year extension.

The notice invites public comment on four specific areas:

1) Whether the proposed collection of information is necessary for the proper performance of DoD functions, including whether the information will have practical utility;
2) The accuracy of DoD's estimate of the burden of the proposed information collection;
3) Ways to enhance the quality, utility, and clarity of the information to be collected;
4) Ways to minimize the burden of the information collection on respondents, including through the use of automated collection techniques or other forms of information technology.

This paper addresses all four areas. Points (iii) and (iv) receive the most extensive treatment, as the authors believe AI-driven automation represents the most consequential near-term opportunity for both the DoD and the contractor community. Point (i) is addressed in Section II; point (ii) in Section XII. The framework developed herein is, however, deliberately broader than this single notice: it is intended as a durable reference for any defense contractor, policymaker, or governance practitioner seeking to integrate AI into defense acquisition compliance.

II. The DFARS Part 244 Information Collection Requirement: Context and Implications

A. Regulatory Architecture of Contractor Purchasing System Oversight

DFARS Part 244 establishes the regulatory framework governing contractor purchasing system reviews (CPSRs). Two provisions are central to the information collection requirement under review:

- **DFARS 244.305—Granting, Withholding, or Withdrawing Approval**: Authorizes administrative contracting officers (ACOs) to evaluate the acceptability of a contractor's purchasing system upon completion of the in-plant portion of a CPSR, to approve or disapprove the system, and to pursue correction of identified material weaknesses. Disapproval triggers a requirement for Government consent to individual subcontracts and may precipitate financial withholding or the exercise of other contractual remedies.
- **DFARS Clause 252.244-7001—Contractor Purchasing System Administration**: Mandates that contractors respond in writing within 30 days of an ACO's initial determination identifying material weaknesses. Following the ACO's evaluation of that response, if material weaknesses persist, the contractor has 45 days to either remediate the weaknesses or submit an acceptable corrective action plan.

The current information collection covers 22 respondents, generating 44 annual responses at an estimated average burden of 8 hours per response, for a total annual burden of 372 hours (OMB Control Number 0704-0253, as reported in the Federal Register notice).

B. The Practical Utility of Purchasing System Oversight

The information collected under DFARS Part 244 is unambiguously necessary for the proper performance of DoD oversight functions. Contractor purchasing systems govern the acquisition of subcontracted goods and services that, in aggregate, constitute a substantial fraction of major defense program costs. Deficiencies in these systems—inadequate competition among subcontractors, insufficient pricing analysis, noncompliance with flow-down requirements—translate directly into cost overruns, schedule delays, and quality degradation in delivered systems. The ACO's ability to approve, condition, or disapprove a contractor's purchasing system is therefore a foundational instrument of defense acquisition stewardship.

However, the *practical utility* of collected information is constrained by the limitations of the current manual, episodic review model. Information gathered at the time of a formal CPSR may be stale by the time corrective actions are assessed. Responses are of variable completeness and analytical depth. And the 30- and 45-day response windows, while appropriate as legal deadlines, may be insufficient for contractors to conduct the root-cause analysis necessary to produce genuinely informative corrective action plans. AI offers a pathway to transform this episodic, compliance-driven data collection into a continuous, analytically rich oversight infrastructure.

C. Current Burden and Its Potential Understatement

The DoD's estimate of 372 annual burden hours for 22 respondents warrants scrutiny. The 8-hour average per response may accurately reflect straightforward cases involving a limited number of discrete material weaknesses. However, for prime contractors operating complex, multi-tier supply chains—where a purchasing system review may implicate hundreds of subcontractor relationships across multiple commodity areas—the actual burden of preparing a substantive, evidence-based response within 30 days is likely to be substantially higher. The estimate also does not appear to capture indirect burden: the time spent by procurement, finance, and legal personnel coordinating data collection across enterprise systems, or the opportunity cost of diverting senior staff from mission-critical activities during a response cycle.

This concern is not merely academic. An underestimated burden figure may cause DoD to underweight the value of investments in automated collection tools that could reduce that burden, and may cause OMB to apply an inappropriately low threshold in evaluating the proportionality of the collection requirement. Section XII returns to this point with specific recommendations.

III. Foundational Principles of Governance in the Defense Sector

Effective corporate governance in the defense industry is distinguished from general commercial governance by the sovereign nature of the primary customer, the classified character of much operational information, and the life-safety consequences of component or system failure. Seven principles are foundational.

Transparency demands open disclosure of material information—board decisions, financial performance, and, critically, quality or safety issues affecting systems destined for military use [7]. In the defense context, transparency extends to supply chains: contractors must be able to disclose the provenance of rare-earth materials, electronic components, and embedded software to prevent counterfeiting, unauthorized substitution, or supply chain infiltration by adversarial actors.

Accountability establishes unambiguous chains of responsibility—from the production floor through the executive suite to the board of directors—and extends upward to oversight bodies including the Defense Contract Audit Agency (DCAA), the Defense Contract Management Agency (DCMA), and congressional oversight committees [8]. For defense contractors, accountability also flows downward through the subcontract hierarchy, as prime contractors bear regulatory responsibility for the purchasing systems and compliance postures of their major subcontractors.

Fairness governs both internal corporate relationships—preventing self-dealing among directors and protecting minority shareholders—and external procurement relationships, ensuring that small businesses and disadvantaged firms receive equitable access to subcontracting opportunities consistent with FAR Part 19 requirements.

Responsibility demands ethical behavior that transcends legal minimums. In defense manufacturing, this principle carries particular moral weight: a component failure in a weapon system or a military aircraft can result in the loss of military personnel, civilian casualties, or mission failure with geopolitical consequences [1]. Responsibility encompasses both positive duties (producing reliable, fully tested systems) and negative duties (refusing to conceal anomalies to preserve contract performance incentives).

Risk management requires the systematic identification, assessment, and mitigation of operational, financial, and compliance risks across a portfolio that may include export-controlled technologies, classified programs, sole-source supply relationships, and long-term fixed-price contracts subject to Nunn-McCurdy breach reporting. Defense contractors face a distinctive risk topology that general enterprise risk frameworks do not adequately address.

Compliance with ITAR, the Export Administration Regulations (EAR), DFARS, and applicable military standards (MIL-STD) forms a non-negotiable governance foundation, enforced through mechanisms including PCAOB oversight of publicly traded contractors and DCAA audit authority [9].

Ethical culture—the shared values, norms, and behavioral expectations that shape decision-making throughout the organization—must be actively and continuously cultivated by senior leadership. The "spiral of silence" that enabled the Enron fraud and the Boeing safety culture failures operated through the same organizational dynamic: individuals who detected problems self-censored because they perceived organizational norms as hostile to dissent [2]. Board-level governance mechanisms must actively disrupt this dynamic through robust whistleblower protections, visible leadership modeling, and genuine psychological safety for quality and safety reporting.

IV. Classical Governance Theories and Their Application to Defense Contracting

Six classical governance theories provide the theoretical infrastructure for the AI governance framework developed in this paper. Each captures a distinct dimension of the principal-agent dynamics, resource dependencies, and institutional pressures that characterize the defense contracting environment.

Agency theory identifies the principal-agent problem as the central challenge of corporate governance: managers (agents) may pursue self-interest at the expense of shareholders (principals) when information asymmetry and misaligned incentives permit [11]. In defense firms, agency risks manifest acutely in cost-plus contracting environments where the incentive to control costs is weak, in the concealment of production defects to avoid bonus clawbacks or contract terminations, and in accounting manipulations that obscure true program cost trajectories. The classified nature of many defense programs compounds information asymmetry, limiting the board's ability to independently verify management representations.

Stewardship theory offers a complementary perspective: executives may be intrinsically motivated to act as stewards of the organization's long-term interests, seeking reputational and mission-oriented rewards rather than short-term personal gain [12]. The mission-driven character of national security work creates favorable conditions for stewardship culture in defense firms—provided that governance structures support, rather than undermine, long-term orientation.

Resource dependency theory conceptualizes the board as a boundary-spanning mechanism that provides access to critical external resources: regulatory approvals, military contracts, political relationships, and strategic supply chain partnerships [12]. For defense contractors, the loss of a single major program or a critical single-source supplier can pose existential risks that dwarf those faced by diversified commercial firms.

Stakeholder theory extends the governance obligation beyond shareholders to encompass employees, local communities, taxpayers who fund the defense budget, and the military personnel whose operational effectiveness and safety depend on the quality of delivered systems [13]. The Department of Defense itself must be understood as a stakeholder with governance interests that differ in character from those of commercial customers.

Institutional theory explains organizational behavior as a response to coercive, normative, and mimetic pressures from the external environment [14], [15]. Defense contractors operate within a dense institutional environment—military standards, ITAR, DFARS, FAR, congressional oversight, and informal professional norms of the defense acquisition community—and must maintain legitimacy across all of these dimensions simultaneously.

Transaction cost economics (TCE) analyzes governance choices—including make-or-buy decisions and monitoring architectures—as functions of asset specificity, uncertainty, and transaction frequency [11]. Defense subcontracting relationships frequently involve highly asset-specific investments (specialized tooling, classified facilities, unique technical expertise) that create lock-in and elevate the cost of switching suppliers, with direct implications for purchasing system design and oversight.

V. Lessons from Corporate Governance Failures: Implications for Defense

The defense industry is not immune to the classes of governance failure that have produced catastrophic outcomes in other sectors. An analysis of three canonical failures—Enron (financial fraud), Kia/Hyundai (systematic product quality concealment), and Boeing 737 MAX (safety-critical software and cultural failure)—reveals structural vulnerabilities that are directly present in defense contracting environments and that AI governance mechanisms are specifically designed to address.

A. Enron (2001): Financial Fraud and the Limits of Audit

The Enron collapse represents the archetype of financial governance failure driven by information manipulation at scale. Executives systematically deployed mark-to-market accounting, special purpose entity (SPE) structures, and deliberate disclosure obfuscation to conceal billions in debt while inflating reported earnings [2]. The failure of Arthur Andersen's audit function, the passivity of the board's audit committee, and the complicity of external legal and financial advisors demonstrated that periodic, human-conducted audits are vulnerable to sophisticated, sustained manipulation. The Sarbanes-Oxley Act of 2002 addressed some of these vulnerabilities through enhanced internal control requirements and auditor independence rules, but the fundamental information asymmetry between managers and overseers remains [8], [9].

For defense contractors, analogous risks manifest in long-duration cost-plus programs where percentage-of-completion accounting can be manipulated, in the use of intracompany transfer pricing to shift costs between government and commercial programs, and in the underreporting of earned value management (EVM) variances to avoid triggering contractual remedies or congressional notification requirements. An AI-driven continuous audit function—analyzing journal entries, related-party transactions, EVM data streams, and revenue recognition patterns in real time—can detect anomalies that periodic audits miss. Natural language processing applied to internal communications can identify patterns of deliberate obfuscation or "tone at the top" that correlates with elevated fraud risk, drawing on the established research utility of the Enron email corpus as a training dataset [2].

B. Kia/Hyundai (2010–2020): Systematic Quality Concealment

The decade-long saga of Theta II engine failures in Kia and Hyundai vehicles—involving engine seizure, uncontrolled fires, and documented pressure to delay defect reporting to regulators—illustrates the governance failure mode of systematic quality concealment under production schedule and financial performance pressure. The parallel for defense manufacturers is direct: propulsion system anomalies in missiles or satellite thrusters, structural defects in aircraft components, or electromagnetic compatibility failures in electronic warfare systems may be discovered during testing but suppressed to maintain contract award schedules or avoid cost-plus adjustments.

AI-powered non-destructive testing (NDT), integrating computer vision analysis of X-ray, ultrasonic, and thermographic imaging data, can detect microscopic material defects—hairline fractures, inclusions, delaminations—that human inspectors miss at production rates [16]. Critically, AI systems can maintain cryptographically integrity-protected logs of all inspection results, preventing the retrospective deletion or alteration of test records that characterized the Kia/Hyundai concealment. Predictive maintenance algorithms applied to production test stand telemetry can identify batch-level failure propensity before components are accepted into finished assemblies.

C. Boeing 737 MAX (2018–2019): Safety-Critical Software and Cultural Failure

The twin crashes of Lion Air Flight 610 and Ethiopian Airlines Flight 302, killing 346 people, resulted from a chain of governance failures that began with a software design decision—the Maneuvering Characteristics Augmentation System's reliance on a single angle-of-attack sensor—and propagated through inadequate safety review, concealment of MCAS functionality from pilots and the FAA, and a corporate culture in which schedule pressure systematically overrode engineering rigor [1]. The Boeing case is particularly salient for defense governance because the failure mode—autonomous software behavior that was not fully understood by operators, combined with management pressure that suppressed safety dissent—is directly replicated in the development of autonomous weapon systems, AI-enabled flight control systems, and cyber-physical defense platforms.

AI-augmented software verification, incorporating formal methods, model checking, and systematic fuzz testing, can enumerate edge cases in safety-critical code that human test teams fail to anticipate. For a defense guidance system, AI can simulate millions of sensor failure scenarios to verify that the system degrades gracefully rather than catastrophically. NLP applied to design review documents, change control board minutes, and engineering email communications can identify patterns of "production pressure" language—schedule commitments displacing test completion requirements—and flag these as cultural leading indicators of safety risk. AI-enforced change control workflows can ensure that no software update to a fielded system is deployed without completing all required safety reviews.

D. Cross-Cutting Synthesis: The Multi-Dimensional Governance Imperative

Defense contractors simultaneously face all three failure modes: financial fraud risks (Enron-type) in complex contract accounting and supply chain pricing; product quality risks (Kia-type) in precision-manufactured components where defect rates must approach zero; and safety-critical software risks (Boeing-type) in the autonomous and semi-autonomous systems that increasingly define modern defense capabilities. An effective AI governance architecture must address all three dimensions within a unified oversight platform, enabling the board to maintain a holistic risk picture rather than managing each failure mode through separate, siloed processes.

VI. Mapping AI Capabilities to Classical Governance Theories

Each classical governance theory identified in Section IV maps to specific AI capabilities that operationalize its prescriptions within the defense context. This mapping constitutes the theoretical core of the proposed framework.

A. Agency Theory: AI as an Information Asymmetry Equalizer

Agency theory prescribes mechanisms that reduce information asymmetry between principals and agents. AI operationalizes this prescription through real-time, auditable data streams that remove managerial control over what information reaches the board. Machine learning algorithms continuously monitor production metrics, financial transactions, purchasing system compliance indicators, and executive decisions, feeding verified data to board dashboards that bypass the management reporting layer. For purchasing system oversight specifically, AI systems analyzing DFARS Part 244 compliance metrics can detect unauthorized subcontract placements, inadequate competition documentation, or missing pricing analyses before they accumulate into reportable material weaknesses—transforming the ACO's oversight function from a periodic audit into a continuous assurance regime [17].

B. Stewardship Theory: AI as a Long-Term Performance Verifier

Stewardship theory is operationalized through governance mechanisms that align executive incentives with long-term organizational health. AI supports stewardship by maintaining objective, tamper-resistant performance records that enable genuine long-term accountability. AI-enabled balanced scorecards can track program performance against multi-year delivery, quality, and cost objectives—providing the board with verified data to evaluate whether executives are acting as stewards of the firm's defense mission or optimizing for short-term compensation metrics [5]. NLP analysis of internal communications can detect when managers are acting in genuine stewardship (proactively surfacing quality issues) versus in self-interested concealment.

C. Resource Dependency Theory: AI as Strategic Intelligence

Resource dependency theory prescribes board-level attention to the external resource environment. AI transforms this from a periodic board agenda item into a continuous intelligence function. Machine learning models analyzing geopolitical news, export control regulatory changes, supplier financial health indicators, and commodity price trends provide early warning of resource disruptions months before they materialize in delivery failures or cost spikes [18]. For DFARS Part 244 compliance specifically, AI monitoring of subcontractor financial health and performance metrics can identify suppliers approaching insolvency or quality deterioration before they trigger purchasing system material weaknesses.

D. Stakeholder Theory: AI as Continuous Compliance Monitor

Stakeholder theory prescribes governance mechanisms that track and respond to the full range of stakeholder interests. AI operationalizes this through continuous, multi-channel compliance and sentiment monitoring. NLP processes regulatory updates from DoD, DCAA, DCMA, and congressional sources; analyzes public reporting and news media for emerging contractor reputation issues; and monitors internal whistleblower channels for early signals of governance deterioration. For defense contractors with significant environmental footprints (propellant manufacturing, testing range operations), AI can monitor regulatory filings and community sentiment simultaneously, enabling proactive stakeholder engagement [19].

E. Institutional Theory: AI as Regulatory Conformity Engine

Institutional theory prescribes mechanisms for maintaining legitimacy across the dense regulatory environment of defense contracting. AI serves as a regulatory conformity engine: trained on the complete corpus of applicable regulations, military standards, and agency guidance documents, the system automatically identifies regulatory changes, maps them to affected internal processes, and generates updated compliance checklists. When new AI ethics guidelines for autonomous weapon systems emerge, or when DFARS clauses are amended, the governance AI immediately propagates the required process updates throughout the organization, reducing the lag between regulatory change and organizational response that creates compliance gaps.

F. Transaction Cost Economics: AI as a Monitoring Cost Reducer

TCE prescribes governance structures that minimize the sum of transaction, monitoring, and enforcement costs. AI dramatically reduces monitoring costs in multi-tier defense supply chains by automating the continuous review of supplier quality data, delivery performance, financial health, and regulatory compliance. For purchasing system oversight under DFARS Part 244, AI can automatically verify that all required pricing analyses, competition documentation, and flow-down clause compliance records are present and complete for each subcontract action—reducing the manual labor required for CPSR preparation and ACO review simultaneously. Smart contract mechanisms can further reduce enforcement costs by automating payment upon verified delivery of conforming items [20].

VII. Preventing the Kill Issue: AI and the Ethics of Safety-Critical Governance

The most severe governance failure in defense manufacturing is the one that directly costs lives: a defective missile guidance component that causes a fratricide event, a satellite thruster failure that leaves a communications satellite uncontrolled, or an autonomous system software fault that results in civilian casualties. This ”kill issue”—the possibility that a governance failure directly produces lethal outcomes—imposes ethical obligations on defense contractor boards that transcend ordinary commercial fiduciary duties.

A. AI-Driven Anomaly Detection for Safety-Critical Components

Modern machine learning algorithms, trained on historical inspection and test data, can identify microscopic anomalies in manufacturing output that human inspectors miss at production scale [16], [21]. Continuous analysis of X-ray inspection images, vibration signatures, material composition reports, and assembly process telemetry enables the automatic quarantine of suspect components before they enter the finished goods inventory. Critically, all inspection events are recorded in integrity-protected logs that cannot be altered retrospectively—eliminating the possibility of the retrospective concealment that characterized the Kia/Hyundai and Boeing cases.

B. Predictive Risk Assessment for Catastrophic Failure Modes

Machine learning models trained on historical test data and informed by physics-based failure mode analysis can predict which production batches carry elevated in-field failure probability. For precision strike systems, even a sub-percent undetected defect rate may carry unacceptable safety consequences. AI-driven predictive analytics enables boards to mandate reinspection or redesign cycles before fielding, embodying the classical ethical principle of precautionary risk management—the obligation to prevent foreseeable harm even under residual uncertainty [4].

C. Operationalizing Classical Ethics Through AI Governance

Classical ethical frameworks, which have historically operated as aspirational guides rather than operational constraints, become genuinely actionable through AI governance systems. **Kantian deontology**, with its categorical imperative requiring that moral rules be universalizable, is operationalized when AI systems audit production decisions against a codified ruleset—flagging any deviation from mandated safety protocols as a deontological violation regardless of schedule or cost consequences [4], [6]. When a production supervisor overrides a required safety test, the AI documents the deviation, alerts the audit committee, and initiates an escalation workflow—making rule violations visible and consequential in real time.

Utilitarian ethics is operationalized through AI's capacity to model the expected consequences of governance decisions at scale. When a contractor faces the temptation to ship a component with a marginal anomaly to avoid a contract penalty, an AI governance system can quantify the probabilistic expected harm—integrating estimated failure probability, operational context, and downstream consequences—against the financial benefit of on-time delivery, enabling a genuinely utility-informed board decision rather than one driven by short-term financial pressure.

D. Explainability as an Ethical Governance Requirement

For defense governance applications, every AI-driven recommendation that affects safety or procurement must be explainable to non-specialist decision-makers. Techniques such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) enable AI systems to generate plain-language justifications for their outputs, supporting the classical governance virtue of accountability: the board can explain and defend its AI-informed decisions to regulators, military customers, and the public [7]. Explainability is not merely a technical requirement—it is an ethical one, ensuring that AI systems do not substitute opaque algorithmic authority for the human judgment that governance accountability requires.

VIII. AI Solutions for DFARS Part 244 Information Collection

This section directly addresses the Federal Register notice's invitation for comment on ways to enhance the quality, utility, and clarity of information collected under DFARS Part 244, and ways to minimize respondent burden through automated collection techniques.

A. Continuous Automated Monitoring as a Structural Alternative to Episodic Review

The most consequential AI contribution to DFARS Part 244 compliance is not incremental improvement to the existing episodic CPSR model but rather its structural transformation. An AI-driven continuous monitoring architecture would deploy machine learning models—trained on CPSR findings, DFARS regulatory text, and contractor purchasing system data—to monitor purchasing system performance on an ongoing basis. Key indicators such as competition rates by commodity, pricing analysis completion rates, flow-down clause compliance, and supplier qualification status would be tracked in real time. Deviations from acceptable performance thresholds would trigger automated alerts, enabling contractors to initiate corrective action *before* anomalies accumulate into the material weaknesses that trigger formal ACO determinations.

This approach addresses the Federal Register's point (iii)—enhancing quality and utility—by ensuring that the information available to ACOs is current and analytically rich, rather than a point-in-time snapshot of a purchasing system's state at the moment of an infrequent review. It also addresses point (iv)—minimizing burden—by enabling contractors to identify and resolve issues continuously rather than mobilizing significant organizational resources to respond to formal determinations under tight statutory deadlines.

B. AI-Assisted Response Generation for Formal Determinations

When an ACO issues an initial determination identifying material weaknesses, the 30-day response window creates significant organizational pressure, particularly for prime contractors with complex, multi-tier supply chains. AI can substantially reduce the time and effort required to produce a compliant, substantive response by:

1) Automatically gathering and consolidating relevant evidence from enterprise resource planning (ERP), procurement, quality management, and contract management systems;
2) Generating structured, clause-by-clause draft responses that address each identified material weakness with supporting data;
3) Identifying and surfacing corrective action precedents from prior CPSR cycles;
4) Tracking response deadlines across multiple simultaneous determination cycles and alerting responsible officials to approaching cutoffs.

For final determinations requiring a 45-day corrective action plan, AI can further assist by generating milestone-based remediation plans, assigning ownership of corrective actions across organizational functions, and providing automated progress tracking that ensures commitments are met within the regulatory window.

C. Enhancing Information Quality Through Standardization and Validation

The Federal Register notice specifically invites comment on ways to enhance the quality, utility, and clarity of collected information. AI enables the following quality enhancements:

1) **Machine-readable standardized reporting format**: A structured, machine-readable format for CPSR responses would enable automated analysis and cross-contractor benchmarking by DoD, substantially enhancing the analytical utility of the information collected. DoD should consider developing a standard data schema for DFARS Part 244 responses, analogous to the XBRL tagging standards adopted for SEC financial filings.
2) **Automated data validation**: AI systems can cross-reference reported information against source system data—verifying that stated competition rates, pricing analyses, and supplier qualification records are consistent with underlying transactional data—reducing the risk of inadvertent error or deliberate misrepresentation.
3) **Completeness verification**: NLP models trained on the regulatory text of DFARS Part 244 and clause 252.244-7001 can verify that all required elements of a response have been addressed before submission, reducing the likelihood of incomplete responses that require additional ACO follow-up.
4) **Predictive analytics for future weakness identification**: AI analysis of historical CPSR findings across the contractor community can identify patterns that predict future material weaknesses, enabling proactive DoD outreach to contractors whose purchasing system indicators suggest emerging risk.

D. Minimizing Respondent Burden: Quantitative Analysis and AI Solutions

The DoD's estimate of 372 annual burden hours across 22 respondents (8 hours per response, 2 responses per respondent) merits careful scrutiny. For contractors with straightforward, single-element material weakness determinations, 8 hours may be a reasonable estimate. However, for prime contractors managing thousands of annual subcontract actions across multiple commodity areas and organizational divisions, the coordination effort required to gather evidence, draft a substantive response, obtain legal review, and submit within 30 days likely substantially exceeds this figure.

AI automation can reduce respondent burden through five mechanisms:

1) **Automated data aggregation**: AI systems integrate data from ERP, procurement, quality, and contract management systems, eliminating the manual data-gathering

phase that accounts for a significant portion of response preparation time.
2) **Intelligent document processing**: NLP extracts relevant compliance evidence from supplier documents, subcontract files, and pricing analyses, reducing the manual review burden for response preparation.
3) **Automated compliance checking**: AI continuously verifies supplier certifications, pricing analysis completion, and flow-down clause compliance against DFARS requirements, maintaining a perpetually current compliance inventory.
4) **Self-service compliance portals**: Web-based portals enabling contractors to submit responses and track corrective action status electronically would reduce administrative overhead and enable DoD to process responses more efficiently.
5) **Proactive issue identification**: By identifying emerging weaknesses before they become material, AI continuous monitoring reduces the frequency of formal determination cycles, directly reducing the total annual burden across the contractor community.

A conservative estimate, based on comparable AI implementation outcomes in financial audit and regulatory compliance contexts, suggests that AI automation could reduce respondent burden by 50–70% within three to five years of implementation [16], [17]. DoD should consider conditioning the three-year extension on a midterm evaluation of AI-assisted compliance tool adoption rates and corresponding burden reduction outcomes.

IX. Opportunities and Benefits of AI-Driven Defense Governance

The strategic case for AI-augmented governance in the defense sector rests on seven categories of benefit that extend well beyond DFARS compliance.

Enhanced transparency and real-time board oversight: AI systems provide continuous monitoring of financial transactions, procurement cycles, and production quality, delivering dynamic risk dashboards to boards of directors that replace retrospective quarterly reports with forward-looking risk intelligence [20]. This shift from periodic to continuous oversight is qualitatively transformative—it enables boards to intervene before governance failures become material events rather than after.

Improved regulatory compliance and audit efficiency: Defense contractors navigating ITAR, EAR, DFARS, FAR, and applicable MIL-STD simultaneously face a compliance burden of exceptional complexity. AI automates compliance checking against this regulatory matrix, reducing human error, lowering internal audit costs, and enabling compliance professionals to focus on high-judgment activities rather than data gathering [16], [21].

Early detection of fraud and misconduct: NLP applied to internal communications, procurement decisions, and financial transactions can identify patterns associated with bribery, false certifications, kickback schemes, and unauthorized technology transfers before they escalate into criminal liability or contract termination events [2].

Strategic supply chain risk management: Machine learning models analyzing geopolitical dynamics, supplier financial health, and commodity market conditions provide early warning of supply chain disruptions that, in defense programs with long lead times and limited supplier alternatives, can cascade into program delays with national security consequences [18].

Strengthened DoD customer relationships: Demonstrating AI-augmented governance and compliance capability to the Department of Defense serves as a meaningful competitive differentiator in source selections that increasingly evaluate contractor management systems as part of past performance and responsibility determinations.

Reduced cost of capital: Institutional investors increasingly incorporate governance quality into cost-of-equity assessments. A demonstrable commitment to AI-enhanced accountability—particularly in a sector where governance failures carry reputational and legal consequences of exceptional magnitude—can lower borrowing costs and attract capital from ESG-oriented funds [22].

Operational efficiency gains: Automating manual audit, compliance, and reporting tasks enables skilled technical and professional personnel to be redirected from administrative compliance activities to mission-critical engineering, program management, and innovation functions [19].

X. Challenges and Barriers to AI Governance Implementation

A rigorous assessment of AI governance opportunities requires equal candor about implementation barriers. Five categories of challenge merit specific attention in the defense context.

Data quality and legacy system integration: Many established defense contractors operate manufacturing and financial systems that were not designed for AI integration. Decades of component quality records, procurement data, and financial ledgers may reside in incompatible formats across disconnected legacy systems. The upfront investment required to cleanse, integrate, and normalize this data is substantial and must be realistically accounted for in any AI governance business case [17].

Algorithmic explainability and procedural fairness: Defense procurement decisions frequently involve the exercise of discretionary judgment by contracting officers and boards of directors who bear legal and ethical accountability for their decisions. Black-box AI recommendations that cannot be explained in terms accessible to non-specialist decision-makers are inconsistent with the procedural fairness requirements of government contracting and the fiduciary duties of corporate directors [7]. Governance AI deployments must prioritize explainable AI (XAI) architectures from the design stage.

Cybersecurity and adversarial AI risk: An AI governance system in a defense contractor environment is itself a high-value intelligence target. A compromised AI system could falsely certify defective components as compliant, generate misleading audit trails, or exfiltrate sensitive program data. Governance AI must be protected through zero-trust network architectures, adversarial robustness testing, and continuous anomaly monitoring of the AI system's own outputs.

Regulatory uncertainty and liability allocation: No specific federal standards currently govern the use of AI in defense contractor governance or DFARS compliance functions. The legal liability implications of erroneous AI audit findings—particularly if a contractor relies on an AI system's assurance of purchasing system compliance and subsequently receives a material weakness determination—are unresolved. DoD and industry stakeholders should collaborate to develop clear standards for AI system certification in governance applications and liability frameworks for AI-assisted compliance decisions.

Workforce capability gaps and cultural resistance: Effective AI governance implementation requires data scientists, AI ethicists, and procurement professionals capable of designing, operating, and critically evaluating AI systems—a skill set that most established defense contractors lack in substantial depth. Moreover, long-tenured procurement and finance managers may perceive AI oversight as a threat to their authority, creating cultural resistance that must be managed through deliberate change leadership [10]. Governance AI should be framed to stakeholders as augmenting, rather than supplanting, professional judgment and expertise.

XI. Strategic Considerations for Defense Contractor Leadership

Six strategic considerations should inform board-level deliberation on AI governance investment in the defense context.

Board composition and AI literacy: Effective oversight of an AI-augmented governance system requires board members who can critically evaluate AI system outputs, understand their limitations, and make informed judgments about algorithmic recommendations. Defense contractor boards should consider the appointment of directors with technology, cybersecurity, or AI governance expertise—a qualification that research identifies as increasingly important for maintaining competitive governance legitimacy [14], [15].

Ethical AI framework and institutional commitment: Defense contractors should adopt a principled AI ethics framework—such as the circular model emphasizing transparency, accountability, fairness, and inclusivity—and institutionalize it through governance policies, training requirements, and audit mechanisms [4]. A "Digital Oath" for AI systems, binding governance AI to explicit "Do No Harm" principles and requiring human review of all recommendations that affect safety or legal compliance, can provide meaningful institutional assurance to DoD customers and investors [6].

Whistleblower protection and the limits of AI monitoring: AI monitoring of procurement and financial systems must be carefully demarcated from employee surveillance. Governance AI that processes internal communications for fraud detection must be designed so that legitimate safety and compliance reporting through protected whistleblower channels is not chilled by the perception of pervasive monitoring. Sarbanes-Oxley whistleblower protections are non-negotiable constraints on governance AI system design [9].

Phased implementation with high-risk area prioritization: Rather than enterprise-wide simultaneous deployment, defense contractors should implement AI governance tools in a phased sequence, beginning with the areas of highest risk and regulatory consequence—such as supplier invoice verification, purchasing system compliance monitoring, and non-destructive testing image analysis—before expanding to lower-risk applications [19]. This approach manages implementation risk, enables organizational learning, and generates early evidence of return on investment.

Liability, insurance, and contract implications: Boards must ensure that governance AI deployment does not inadvertently create unmanaged legal exposure. If an AI system fails to detect a defect that subsequently causes program failure or personal injury, the allocation of liability between the contractor, the AI system vendor, and the DoD customer requires advance legal analysis. Directors' and officers' insurance policies should be reviewed to confirm coverage for AI-assisted governance decisions.

Competitive positioning and first-mover considerations: The defense contractor that establishes a credible, DoD-validated AI governance capability may realize significant competitive advantages in source selections that evaluate contractor management systems. However, first-mover costs in uncharted regulatory territory are real, and leadership must weigh these against the long-term strategic value of governance leadership [20].

XII. Structured Recommendations: Formal Response to OMB Control Number 0704-0253

The following recommendations are submitted in direct response to the four areas of public comment invited by the Federal Register notice regarding OMB Control Number 0704-0253.

A. Area (i): Necessity and Practical Utility of the Information Collection

The information collection under DFARS Part 244 is necessary for the proper performance of DoD oversight functions and should be extended. Contractor purchasing systems govern subcontracted expenditures that represent a substantial fraction of major defense program costs, and ACO oversight of these systems is a foundational instrument of defense acquisition stewardship. However, the practical utility of the information collected can be substantially enhanced through AI-enabled continuous monitoring, standardized machine-readable reporting, and predictive analytics. The authors recommend that DoD condition the three-year extension on a structured evaluation of AI-assisted compliance tool adoption and its impact on information quality and practical utility.

B. Area (ii): Accuracy of Burden Estimates

The DoD's estimate of 372 annual burden hours (22 respondents, 2 responses each, 8 hours per response) is plausible for straightforward cases but is likely to understate the true burden for prime contractors managing complex, multi-tier supply chains. The authors recommend that DoD conduct a supplemental burden analysis stratified by contractor size and purchasing system complexity, and that this analysis be used to refine burden estimates in the three-year extension period. Simultaneously, DoD should establish baseline metrics for current burden levels to enable rigorous evaluation of AI-driven burden reduction over the extension period.

C. Area (iii): Enhancing Quality, Utility, and Clarity of Information

The authors recommend the following specific enhancements:

1) Develop a standardized, machine-readable data schema for DFARS Part 244 responses, modeled on XBRL financial reporting standards, to enable automated analysis and cross-contractor performance benchmarking;
2) Establish DoD-approved AI validation tools that contractors can use to verify the completeness and consistency of their responses prior to submission;
3) Create a centralized DoD analytics capability to extract systemic purchasing system risk patterns from the information collected, enhancing its value as a defense acquisition intelligence resource;
4) Implement a pilot program to evaluate AI-driven continuous monitoring as an adjunct to the existing episodic CPSR model, selecting 5–10 volunteer prime contractors across diverse defense industry segments to test and validate the approach before broader implementation.

D. Area (iv): Minimizing Respondent Burden Through Automation

The authors recommend the following:

1) DoD should issue guidance endorsing the use of AI-assisted tools for DFARS Part 244 compliance monitoring and response preparation, providing clarity on the permissibility of AI-generated draft responses subject to human review and certification;
2) DoD should develop and make available a reference architecture for AI-enabled purchasing system monitoring, enabling smaller contractors without in-house AI capability to adopt standardized solutions;
3) The OMB burden estimate for the three-year extension period should incorporate a phased reduction trajectory reflecting projected AI automation adoption, with target burden reductions of 30% by year two and 50–60% by year three;
4) DoD should establish electronic self-service portals for CPSR response submission and corrective action tracking, replacing paper-based processes that add administrative overhead without enhancing information quality.

XIII. Conclusion and Policy Implications

This paper has developed a comprehensive, theory-grounded framework for AI-enhanced corporate governance in the U.S. defense industrial base, with direct application to the information collection requirements under DFARS Part 244. The foundational principles of transparency, accountability, fairness, responsibility, risk management, compliance, and ethical culture remain the bedrock of sound defense contractor governance. The contribution of this framework is to demonstrate, with theoretical rigor and practical specificity, how AI can operationalize these principles at a scale and speed that human governance processes alone cannot match.

The mapping of AI capabilities to classical governance theories reveals a coherent logic: agency theory's prescription for reduced information asymmetry is fulfilled by AI's continuous audit function; stewardship theory's orientation toward long-term value is supported by AI's verified performance metrics; resource dependency theory's emphasis on strategic external intelligence is operationalized by AI's supply chain monitoring capability; stakeholder theory's multi-constituency accountability is served by AI's continuous compliance monitoring; institutional theory's demand for regulatory legitimacy is addressed by AI's conformity engine; and transaction cost economics' efficiency prescription is met by AI's dramatic reduction in monitoring costs.

The lessons of Enron, Kia/Hyundai, and Boeing 737 MAX are not merely historical cautionary tales. They are structural failure patterns that are directly replicable in defense contracting environments—and that AI governance mechanisms are specifically designed to detect and prevent. The ”kill issue” in defense manufacturing demands governance mechanisms that can achieve near-zero-defect assurance at production scale, and AI-augmented inspection, testing, and quality management systems represent the most credible available pathway to that standard.

For policymakers, this analysis suggests five priority actions:

1) Develop specific AI governance standards for defense contractors under DFARS, including certification requirements for AI systems used in safety-critical manufacturing and compliance functions;
2) Establish a DoD-sponsored pilot program for AI-driven purchasing system monitoring under DFARS Part 244, using the OMB Control Number 0704-0253 extension period as the implementation window;
3) Fund research on explainable AI architectures specifically designed for defense audit and compliance applications, ensuring that AI governance tools are compatible with the procedural fairness requirements of government contracting;
4) Create protected information-sharing mechanisms for AI governance best practices across the defense industrial base, enabling smaller contractors to benefit from prime contractor implementation experience;

5) Develop clear regulatory guidance on the liability implications of AI-assisted compliance decisions, providing the legal clarity that contractor boards require to authorize AI governance investments.

For defense contractor boards and senior leadership, the path forward requires: deliberate AI literacy development at the board level; phased implementation beginning with highest-risk governance domains; robust cybersecurity protection of AI governance systems; and explicit commitment to the ethical AI frameworks that align AI governance tools with the broader obligations of responsible defense industrial citizenship.

The central thesis stands: AI is not a governance panacea, and it cannot substitute for the human judgment, institutional integrity, and ethical commitment that effective defense contractor governance requires. It is, however, a powerful and increasingly necessary meta-governance layer. When embedded within robust traditional governance structures, calibrated to the specific regulatory and ethical obligations of the defense sector, and guided by the classical governance theories that have defined the discipline, AI can help defense contractors fulfill their role as reliable, accountable, and ethically grounded partners in U.S. national security—while simultaneously reducing the administrative burden that regulatory compliance imposes on the contractor community and the Government alike.

Formal Comment Submission Statement

This paper is submitted as a formal public comment in response to Federal Register Volume 91, Issue 103 (May 29, 2026), Federal Register Document No. 2026-10731, regarding the proposed extension of information collection requirement OMB Control Number 0704-0253, Defense Federal Acquisition Regulation Supplement, Subcontracting Policies and Procedures, Docket Number DARS-2026-0202. The analysis and recommendations herein address all four areas of public comment invited by the Defense Acquisition Regulations System. Comments may also be submitted electronically at https://www.regulations.gov or by email to OUSD-AT-L-DARS-info@mail.mil, referencing OMB Control Number 0704-0253, by the July 28, 2026 comment deadline.

Declaration

The views expressed are those of the author and do not represent any affiliated institution. This work is conducted as independent research. This paper is a conceptual and review contribution; all proposals and findings are derived from and grounded in the cited literature. The author does not claim novel experimental results but advances an integrated framework synthesizing existing research for application to defense governance and regulatory policy.

Portions of this manuscript were drafted with the assistance of AI writing tools to improve clarity and organization. All AI-generated content was reviewed, edited, and verified by the author for coherence and factual accuracy; however, readers are encouraged to independently verify cited claims. The LaTeX code was developed with the assistance of AI tools.

References

[1] R. Desai, "The Boeing 737 MAX Crisis: Public Trust, Aviation Safety, and Corporate Accountability?" pp. 463–469. [Online]. Available: https://ajosr.org/wp-content/uploads/journal/published_paper/volume-3/issue-6/ajsr2025_nFGEQxgg.pdf

[2] K. Taylor-Neu, A. S. Rahaman, G. D. Saxton, and D. Neu, "Tone at the top, corporate irresponsibility and the Enron emails," vol. 37, no. 9, pp. 336–364. [Online]. Available: https://doi.org/10.1108/AAAJ-12-2023-6792

[3] J.-E. Pozner, A. Mohliver, and C. Moore, "Shine a Light: How Firm Responses to Announcing Earnings Restatements Changed After Sarbanes-Oxley," vol. 160, no. 2, pp. 427–443. [Online]. Available: https://www.jstor.org/stable/45278228

[4] S. A. Hosseini Tabaghdehi and O. Ayaz, "AI ethics in action a circular model for transparency, accountability and inclusivity." [Online]. Available: https://doi.org/10.1108/JMP-03-2024-0177

[5] P. K., G. Sehgal, U. M. Padwal, D. Shah, D. Tewari, and S. Rana, "Corporate Governance and Ethical Management Practices: Building Trust in Modern Enterprises." vol. 2, no. 4, pp. 5388–5396. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=935212d8-a2e3-35c6-9b0b-c492e152e483

[6] M. Najm and p. u. family=Amaral Giovani, given=Joao Francisco, "ICAIMT — The Digital Oath for Corporate Governance: A First Model." vol. 24, no. 6, pp. 1–17. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=799ff2b9-d658-35fa-ae6c-8b389b2ffacb

[7] C. Li, "AI-Driven Governance: Enhancing Transparency and Accountability in Public Administration," vol. 1, no. 1, pp. 1–16. [Online]. Available: https://www.dsvg.cc/dsvg/article/view/1

[8] M. W. PEREGRINE, "The Continuing Governance Relevance of Sarbanes-Oxley: Two decades on, Sarbanes-Oxley remains a touchstone for ethics, oversight and the modern director's duty." vol. 50, no. 2, pp. 58–61. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=b75aab09-47c1-3509-b504-67e0a3b825ca

[9] D. L. King and C. J. Case, "SARBANES-OXLEY ACT AND THE PUBLIC COMPANY ACCOUNTING OVERSIGHT BOARD'S FIRST ELEVEN YEARS." vol. 7, no. 1, pp. 11–22. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=ef4fb984-b875-33a6-bc0b-630ad8e62abd

[10] I. Gallego-Álvarez and L. Rodriguez-Dominguez, "The Role of Organisational Factors and Corporate Governance in Business Ethics Practices." vol. 8, no. 1, pp. 1–17. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=8223f3c1-b335-39a4-b094-57b06d8d3adc

[11] C. D. Cole and R. Schneider, "Corporate Governance and Capital Structure: Extending Agency Theory From Executives to the Board of Directors." vol. 20, no. 6, pp. 108–124. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=617bf704-7fd1-317f-be8e-f88a67a878b3

[12] P. M. Madhani, "Diverse Roles of Corporate Board: A Review of Various Corporate Governance Theories." vol. 16, no. 2, pp. 7–28. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=3e2f85c2-14db-36c5-9716-9a07d48c4fc8

[13] J. Amis, J. Barney, J. T. Mahoney, and H. Wang, "From the Editors—Why We Need a Theory of Stakeholder Governance—And Why This is a Hard Problem," vol. 45, no. 3, pp. 499–503. [Online]. Available: https://journals.aom.org/doi/10.5465/amr.2020.0181

[14] M. A. Witt, S. Fainshmidt, and R. V. Aguilera, "Our Board, Our Rules: Nonconformity to Global Corporate Governance Norms," vol. 67, no. 1, pp. 131–166. [Online]. Available: https://doi.org/10.1177/00018392211022726

[15] ——, "Our Board, Our Rules: Nonconformity to Global Corporate Governance Norms." vol. 67, no. 1, pp. 131–166. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=dd5c7ebc-73e3-3217-a05e-a5672f512b0a

[16] Olubusola Odeyemi, Kehinde Feranmi Awonuga, Noluthando Zamanjomane Mhlongo, Ndubuisi Leonard Ndubuisi, Funmilola Olatundun Olatoye, and Andrew Ifesinachi Daraojimba, "The role of AI in transforming auditing practices: A global perspective review," vol. 21, no. 2, pp. 359–370. [Online]. Available: https://wjarr.com/node/4438

[17] H. H. Samara, H. A. Qudah, H. J. Mohsin, S. Abualhijad, L. Y. Bani Hani, S. Al Rahamneh, and M. Z. AlQudah, "Artificial intelligence and machine learning in corporate governance: A

bibliometric analysis," vol. 44, no. 2, pp. 349–375. [Online]. Available: https://journals.sagepub.com/action/showAbstract

[18] M. Becht, J. Mejino-López, and G. B. Wolff, "Who Controls the Defence Industry?" [Online]. Available: https://www.jstor.org/stable/resrep79335

[19] A. K. Ulfah, K. M. Shavkatovna, M. N. Nuriddinovna, O. L. Erkinovna, R. M. A. O'g'li, and A. Hidayati, "Transformation of good corporate governance through AI: Innovations in enhancing supervision and accountability." vol. 3393, no. 1, pp. 1–21. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=ea17a16a-7cf1-35a4-a95e-04fa1b1e1bf5

[20] O. S. Shaban and A. Omoush, "AI-Driven Financial Transparency and Corporate Governance: Enhancing Accounting Practices with Evidence from Jordan," vol. 17, no. 9. [Online]. Available: https://www.mdpi.com/2071-1050/17/9/3818

[21] F. A. Wassie and L. P. Lakatos, "Artificial intelligence and the future of the internal audit function," vol. 11, no. 1, p. 386. [Online]. Available: https://www.nature.com/articles/s41599-024-02905-w

[22] S. Bamel, Dipika, and S. Singh, "Mapping the research landscape of business ethics, corporate governance, and corporate social responsibility: A bibliometric and citation analysis." vol. 29, no. 1, pp. 31–57. [Online]. Available: https://research.ebsco.com/linkprocessor/plink?id=e83e58ab-f59c-3583-9409-2b5bb8c43f8d